106

An Interview with

LEONARD KLEINROCK

OH 190

Conducted by Judy O'Neill

on

3 April 1990

Los Angeles, CA

Charles Babbage Institute
The Center for the History of Information Processing
University of Minnesota, Minneapolis

nicely. We were the first node, and then in the next three months three more nodes were added - we had a four-node network by December, the famous four-node network. Here is the book I was talking about earlier [*Communication Nets*]; it eventually went into paperback. In this other book of mine [*Queuing Systems, Volume II: Computer Applications*], we see the network that was initially implemented, namely, that four-node network [p. 306]. And then it began to grow; by the middle of 1970 it looked like this. Early in 1971 it began to grow again. And all the while we were making measurements. And they were not just passive measurements - we did more than observe what was going; we were experimenting and generating artificial traffic. We became the network measurement center. Meanwhile my research quickly moved into networks, which I was waiting to do. I began publishing a great deal on the subject.

O'NEILL: Okay. I have some questions on some of the things that you have mentioned. I have a pretty good overview up to the point of your early ARPANET work. You talked about your advisor, Edward Arthur, who was working on a classified project that you didn't have access to for one of the services. Do you know if he took your work and applied it to the network?

KLEINROCK: I don't know, and I sort of doubt it, because he wasn't very heavily coupled into that. I think he attended a conference in which these problems came up, and I don't recall him being heavily coupled. I do remember that there was a thing called UNICOM, I believe, that became interesting in the early to mid 1960s. And I was at a conference in the early 1960s where that was mentioned, and my ideas were made known there. I never worked on that project.

O'NEILL: When you decided to go to UCLA, did you have any interest in or knowledge of the DARPA projects that the Western Data Processing Center had gotten started?

KLEINROCK: No, but that's where Ivan enters the picture. I believe it was 1965, or so, that he came by with his ARPA hat and suggested that we network the three IBM 7090s on campus. There was one in the campus

computing facility, one at Western Data Processing Center, and one over in the medical center. In 1967, Larry entered the picture and began talking about his plans to fund a network implementation. In April 1969 I received an ARPA research contract. It was obvious I was going to take a key role in the forthcoming ARPANET. So I began as principal investigator in April 1969. We still have a major contract with DARPA.

O'NEILL: There is a Program Plan in the DARPA files for the Western Data Processing Center, but it is dated 1963.

KLEINROCK: That was a separate thing. They were providing computing services to some set of entities.

O'NEILL: When you were at Lincoln, or I guess it would have been after you had already come out here, Larry Roberts and Tom Marill did an experiment connecting the TX2 and the Q32. Were you aware of that work?

KLEINROCK: Oh, very much so. Larry spoke to me quite a bit about the problems they were having and the way he was trying to solve it. I was aware of but not participating in that experiment. I know that is what triggered Larry to try to get a system that would work more smoothly. The interesting thing is, as I recall, that part of the motivation for this network is the fact that in 1967, in the mid 1960s DARPA was heavily supporting a lot of people doing work on time-sharing. And every time an investigator got a new contract, the first thing he wanted was a computer - the best and biggest. Pretty soon Larry said, "This is getting ridiculous," because each facility they created evolved into a specialized kind of facility, like the graphics capability at Utah, the database capability at SRI, and the simulation capability at UCLA. So Larry came up with the concept of a resource sharing network, where there would be specialized sites, and if you wanted that special capability, you connect to that site to get it, or you would pull back data or programs and use them locally. That was one of his motivating reasons, namely, to reduce the number of time-sharing systems he had to support.

O'NEILL: Were you also keeping track of what Paul Baran was doing in his Rand reports?

KLEINROCK: I was well aware of his results. In fact I quoted his results in my own dissertation. He had some of the important ideas of packet switching as a system. The thing that really drove my own research was the idea of a message switching network, which was a precursor to the packet switching networks. The mathematical tool that had been developed in queueing theory, namely queueing networks, matched perfectly the model of computer networks. Actually, it didn't match perfectly and I had to adjust that model to fit the realities of computer networks. Then I developed some design procedures as well for optimal capacity assignment, routing procedures and topology design.

O'NEILL: I'm trying to pull it all together. It is clear Larry Roberts had a big part in this, but it is important to understand historically where the roots are and how it came together.

KLEINROCK: Well, you know, Donald Davies of the National Physical Laboratory in the UK was also a key contributor. Donald Davies had the same idea of a packet network independent of Larry Roberts. Both of them post-dated Paul's basic ideas. Now, Paul didn't have the total picture in mind, but he had the essentials there. Unfortunately, Paul's work was not properly recognized in the early literature. His work was in 1964, I believe.

O'NEILL: The reports were actually published in 1964, but they had been written before that and they had had some exposure before that.

KLEINROCK: And I was aware of his work, but again, it didn't play a major part in my thinking until some years later. But I would credit him with the first ideas. Larry and Donald Davies worked in parallel and independently. Unfortunately Donald Davies at NPL never got the proper funding to really let that system blossom physically. I believe he made a one-node switch at NPL. He had the ideas though. He was definitely a forward thinker. Larry fortunately had the resources to make networking happen.

O'NEILL: You said that Larry brought you and others out to Washington to start talking about the specifications for the network. Was that after the principal investigators meeting where he had talked to the host community? I'm assuming you weren't a principal investigator because you said you hadn't gotten involved until later.

KLEINROCK: I attended, but later than 1967. I attended one in late 1968 or early 1969. The meeting was in, of all places, Alta, Utah - a wonderful ski resort.

O'NEILL: So at some point in 1967, you don't recall just when, you went to a meeting in Washington.

KLEINROCK: I remember some of the people there. I remember that Tom Cheatham was there. I believe Herb Baskin was there also, and I think it was he, or it may have been Doug Engelbart, who was hot on time-sharing and said, "If this network can't give me a half a second response time, then I can't use it for time-sharing." So we specified that there should be a half a second response time. We soon found out it could give two tenths of a second response time using the 56 KBPS lines we had. But it was there that it was decided. And I banged my fist on the table and said, "We've got to put measurement software in there. If this is going to be an experiment, we must make measurements. So it evolved; we decided on a reliability criterion, and specified the format and the nature of this network.

O'NEILL: What was the total size of this group that was together in Washington?

KLEINROCK: Ten... on that order.

O'NEILL: Okay. Was this an ongoing series of meetings, or was this just a couple of days?

KLEINROCK: One or two days, I can't remember which. Just one shot.

O'NEILL: Had there been a document prepared for that meeting? Were you working from a draft proposal?

KLEINROCK: No. What happened thereafter I had no access to, but a formal RFQ spec was prepared.

O'NEILL: Okay, so was that your only direct involvement prior to actually getting the IMP installed?

KLEINROCK: No, I believe I was asked to help evaluate some of the proposals as well, but in an unofficial way.

O'NEILL: So you weren't part of a committee evaluating the proposals?

KLEINROCK: No. I was aware of the numbers that were coming in. I was aware that many of them suggested the same computer for the switch, namely, the Honeywell DDP-516.

O'NEILL: In the proposals that you happened to see, was there a lot of variation in how they proposed to do it?

KLEINROCK: I only remember two of them in some detail. I don't remember how much software they specified. And I don't remember, frankly, when the concept of distributed control routing was described in detail. We suggested it, as opposed to centralized control.

O'NEILL: We, being the group in Washington?

KLEINROCK: Yes. It was clear that we wanted no central location which if it failed, would take down the network. We distributed control throughout. But that is the concept; now the particular algorithm I believe was developed at BBN. You may or may not know that BBN was extremely protective of the software. They wrote

the software according to the gross spec that we gave. And they also gave us the network algorithms, but they wouldn't let us look at the code. Now there we were at UCLA; it was our job to test and try to break the network - to experiment with it. And when we looked at the algorithm, we told BBN that there were some problems. There would be loops in the routing procedure and some deadlocks in the flow control algorithm. And they were very reluctant to make changes. They saw us really as a nuisance, for good reason. They had a job to do and we were telling them to do it differently. But the interesting thing is they wouldn't let us see the source code. Even though we suspected a problem in their implementation, we could never see it. We told them, "Here is a problem, fix it." They said it would take six months to fix, so they never did. Finally there was pressure put on them to make the software public (I forget where that pressure came from - it was either DARPA or an external community that said, "Look, the government paid for the software; you don't own it.") Finally they were forced to open up the software. Then we had access to the code, and what we did then was to find the trouble and not only did we tell them to fix it, but we also told them how to fix it. It still took six months. (Laugh)

O'NEILL: When the first IMP was installed here at UCLA, were you personally involved in getting it up and running?

KLEINROCK: Oh, very much so. It cost me a great deal in fact. I was involved with a start-up company at the time, and my participation was severely diminished as a result of that.

O'NEILL: Was that the Linkabit Corporation?

KLEINROCK: Yes. I have since formed another company on my own.

O'NEILL: The Technology Transfer Institute. In fact, I was curious about that name. It kind of invites a question as to what technology, and who is it transferred to?

KLEINROCK: Okay. The concept of that name is that we were going to transfer technology via seminars, that is, to get the word out there quickly. If you want to get to that now, we can. Or maybe you want to wait.

O'NEILL: Why don't we just go ahead, as long as we're talking about it.

KLEINROCK: I published a two-volume book in the mid 1970s. The first was on queueing theory. The second one was on computer applications. It was the second book that talked about computer networks. There were two major chapters on it. It included much of the material here [in the 1964 monograph] and a lot of the work I had done since. This came out in 1976, so that's seven years after the ARPANET began. There was a lot about the ARPANET technology and its development in the intervening years. The book was new, the topic was hot, and my name was known. I thought, "Why not do a seminar to get this information out quickly." Everybody was pushing me to do that. So in the summer of 1976, my wife and I formed a little company, created a "Computer Network" seminar, publicized the seminar, and ran the same seminar in Dallas, Washington, and L.A. It was a big success. So I formed this company and said, "Let's go." And what was the reason for doing this? It's that I had access to the best talent in the performance evaluation and networking area in the country. People like Larry Roberts, like Howie Frank, like Norm Abramson, like Peter Denning, and myself. Using them as lecturers, I created five seminars, five separate seminars, and we launched that in the winter of 1977. We repeated that twice a year; then it went to six seminars twice a year. And it was going along fine, very nicely. Then just to continue the history now, in the spring of 1979, we went from six twice a year to twelve twice a year. And more significantly, we took on James Martin. Do you know James Martin's name?

O'NEILL: Yes.

KLEINROCK: Okay. He had been at IBM. In 1977 he took a sabbatical, got some friends of his in the U.K., and created a five-day seminar; they administered and he spoke - around the world (fourteen places in the world). It was a big success. In 1978 they did it again, and these people in the U.K. were getting tired; so they

decided to parcel out the world. My company was thriving and Jim and I had known each other very well since 1970; we had done some seminars in Israel and Europe together. TTI took over the James Martin seminar in North America. And it has been so ever since. The number of seminars has grown. I was the president of the company at the time; I since got someone to take over that role. I don't handle the day-to-day operations. The things which made it work are as follows. First of all I had the network of the best talent. I knew who the best people were. And I could evaluate them. They had to satisfy two of three criteria, every speaker: they either have written the key text in the field, and be the pioneer in the field, and be an excellent speaker - any two of those three criteria. The second key was the synergistic relationship between my teaching, my research, and the things that I was teaching out in the seminars. It was all the same kind of thing, so it was low overhead to me to create these seminars. It was what I was doing anyway. That's one of the reasons it worked. It is the reason I didn't have to leave the university to make it happen. I got the right management in, and I controlled the quality. And it is still going great.

O'NEILL: When you first started this, was there any concern by, or connection at all, with DARPA, the fact that they were providing money to your networking research?

KLEINROCK. No. In fact, they were very much interested in and happy to have spinoff companies, and to get that information out there.

O'NEILL: Was there anything you would consider active encouragement of that?

KLEINROCK: No. But a variety of the people at DARPA taught for me. Like Larry, like Bob Kahn, like Vint Cerf. They were people at the office at various times. And many of the people they were supporting were teaching for me - Dave Farber, Howie Frank, Carl Sunshine, John McQuillan, etc. But there was no active support, no.

TAPE 1/SIDE 2

O'NEILL: Let's talk a bit more about the situation here at UCLA. I had gotten the impression that there were a lot of students coming in working on the network, doing research on the network. Was that something that you actively set up, tried to work it that way?

KLEINROCK: Very much so. And it had to do with the fact that I had this large ARPA contract that was supporting the things I was interested in, but also helped point the direction as to where the research should go. For example, there were phases of research that I passed through. I started doing some queueing theory, some time-shared modeling, and then moved into networking with support from DARPA. The ARPANET was the first piece of that. There were various waves of research that occurred from there on. Wave one was the ARPANET, for me. Now, of course, the time-sharing work was already being supported by DARPA elsewhere. I was not being supported at the time, but I was heavily involved in that research. All of the money I received was used to support Ph.D. students. But also a big piece of the money that I was given by DARPA was used to support the implementation. You see, UCLA was in charge of creating the host-to-host software. Steve Crocker, Vint Cerf and Jon Postel all worked for me at the time. I had at one time about forty people - secretaries, programmers, managers, faculty and students. It was amazing, I started this huge group. And Steve Crocker was in charge of the software side of it. I had other people in charge of the physical side of the facility, our host, our IMP, and all the rest. But Steve Crocker played a very important role on the software side. Steve was getting his Ph.D.; so was Vint and so was Jon. I filled my students full of networking ideas. Now, as I say, we became the ARPANET measurement center as well. Around that same time DARPA decided to create a transatlantic link to Europe via satellite.

O'NEILL: This was 1972?

KLEINROCK: 1972, 1973 in that range, yes. And we decided to experiment with what is called a broadcast satellite channel using ALOHA as the access scheme. So we began to study packet satellite switching. And some of my students started to move into that research area. For example, one of my students was Mario Gerla who later worked for Howie Frank. We hired him back from Network Analysis Corporation. I'd say Mario's name was the most significant one that came out of the ARPANET research. Then came the satellite stuff. Simon Lam is a name you may or may not know. He is now at the University of Texas. His thesis was studying ALOHA as an access scheme mathematically. And that came out of the satellite switching meetings that we had. There was a set of meetings that went on to develop this technology. I brought Simon with me to these meetings to act as my right-hand man. Out of that came his dissertation. Then ARPA started moving to packet radio on a metropolitan area networking basis. The application was to a soldier running around the field with digital radios, or possibly tanks moving across a battlefield; that led to the SURAN, survival radio network, and the whole packet radio project began. We started studying CSMA, carrier sense multiple access, which eventually led to ETHERNET. Fouad Tobagi was also my student; he's now at Stanford. For his thesis, I told him "You're going to evaluate CSMA." He did a beautiful job. There is a sequence there: Gerla, Lam, and Tobagi. All at universities, all well known, all right now. Following that came some work on local area networks that DARPA didn't really promote heavily. But it was a technology that came out of the packet radio project. The same ideas for sharing a ground radio network are the same ideas you need to share a bus or a ring-type communications medium. Bob Metcalf, who was present in all of this (he was up at Xerox at the time) came up with the idea of CSMA/CD, with collision detect to use on a wire, and out of that, of course, came Ethernet and all the rest followed. Some other of my students' names may interest you. Bill Naylor who was one of the programmers under Steve Crocker, did a dissertation for me. His job was to do the measurement of the ARPANET. He and a fellow named Gerry Cole were doing measurements theses. They took data and they created and analyzed models. Out of the packet radio came a young man named John Silvester. He is well known in the packet radio research field. He is over at USC right now. And there are many others. Do you care about these names?

we did the ARPANET measurement. Following that we were tasked with doing some satellite measurement and some packet radio measurement. I spun off the packet radio measurement as well. So there were two spawns that took place out of the original ARPA support.

O'NEILL: Were you free to make these subcontracts?

KLEINROCK: No, no, no. I did it formally through the ARPA office. I said, "Look you should be supporting these guys on their own."

O'NEILL: And did that take away from the money you were getting?

KLEINROCK: Well, my pile was growing. It was growing bigger than I wanted. It was approaching a million dollars a year. I don't need that much money. As it is now, I'm being supported at about two thirds of a million, and that is probably going to stop come next renewal, because they can't afford so much money for "theory and performance evaluation," which is unfortunate. I'm going to shrink considerably this next round. Right now I'm supporting seven Ph.D. students and a couple of secretaries. That will go way down. A lot of it has to do with new structure in the ARPA office, which has totally changed.

O'NEILL: That brings up another point, the changes in the ARPA office over time. Obviously you have seen that change over several years now. Maybe we should back up, and why don't you describe what your interaction was like with the office initially.

KLEINROCK: Okay. The reason I got the contract initially was because Larry and I knew each other well, we respected each other, and the ARPANET thing got going. At that time he was not the head of the office, Bob Taylor was. Then Bob left and shortly thereafter Larry took over. Things with DARPA were beautiful at that time. After Larry left, Bob Kahn eventually moved in as the head of IPTO.

O'NEILL: Licklider came back for a short time, and then there was Russell, and then Kahn.

KLEINROCK: True. Kahn had a very long tenure, and Bob and I got along famously. In fact, Bob was in the initial ARPANET experiment, as you well know. He came out here as a BBN rep, and did some of the initial measurements with me. We had excellent rapport. Eventually, Bob Kahn and Bob Cooper had some problems regarding support of networking.

O'NEILL: Who is Bob Cooper?

KLEINROCK: Bob Cooper was the head of all of DARPA toward the end of Bob Kahn's tenure.

O'NEILL: Where do you look now for funding?

KLEINROCK: Me personally? I'll probably not look for large contracts. I think I'll be happy with a smaller amount and fewer students. It is a big administrative drain to run a large contract. I've been running six, seven, eight students continually, and I may just shrink down to three or four, reduce my staff and just manage with the money I have. Somehow, it is probably time to cut it down.

O'NEILL: But that would still be ARPA money then that you would run on?

KLEINROCK: Yes, to the ARPA money. I might also go to NSF or I might go to IBM, if I feel I need some more, but I probably won't.

O'NEILL: If you had wanted to maintain your level, are NSF and IBM the places you would go?

KLEINROCK: Possibly.

O'NEILL: I would like to back up a little bit, again, to the working relationship when the ARPANET was being installed. You mentioned your inter-relationship with BBN. Can you talk more about that? Not only BBN but obviously the Network Analysis Corporation, and SRI perhaps.

KLEINROCK: Okay. Let me talk about NAC first. That goes back to this picture again (in Volume II). In the early days there was no problem in designing the topology of the ARPA network. I mean, how else would you connect four nodes? And maybe the same for ten nodes where this group is on the east coast and this group is on the west coast - it is pretty obvious. As we got into larger networks, topology became a major issue. I had known Howie Frank for many years prior to the whole ARPA thing. In fact I met him in 1966 when I did a lecture for him at Berkeley - he was on the faculty at Berkeley at the time - and so was Ivan Frisch. They had written this very important book on network flow. In fact, I almost got killed going to lecture for them up in Berkeley at the time. I was coming out of Sequoia in a station wagon and the brakes failed going downhill at forty-five miles an hour... but, that's another story. (Laugh) Actually, in 1968, Howard Frank and myself, and a bunch of other guys, including Ivan Frisch, went to Washington, D.C. to work for the Executive Office of the President.

O'NEILL: So you were part of the same group that Howard Frank was involved with?

KLEINROCK: Half and half. Do you know about that whole effort? A guy named David Rosenbaum was the one who promoted this thing, and he knew Howie. Basically, he got the government to agree to create this group, that summer, to try to show how network flow theory could be useful to government problems. I was not on the full-time team. I spent two weeks there. The problem they chose to work on was a problem for the Federal Power Commission on designing natural gas pipeline networks. Great work! Very important algorithms came out! And the most important person there, to be honest with you, is none of the names I mentioned - was a guy named Danny Kleitman, who is a mathematics professor now at MIT. Brilliant guy and he came up with the key algorithm. We made some important progress. We saved the government $20,000,000 dollars over a

few years, with a small investment. Shortly thereafter Howie, Ivan and Dave formed the Network Analysis Corporation.

I went there to lecture a couple of times. But I knew the whole group. In about 1971 or so Larry was at my house, and I suggested that he meet Howie Frank to assist in the topology design problem. So I put them together. And it was a click. Then Larry gave NAC the contract for doing the topological design for the ARPA network. That's the history on that, and it's continued ever since. UCLA was the first node, and SRI was the second node. So this was the first link in the ARPA network - there it is, there it is, and there it isn't. Around this time we were busy simulating the network and also measuring it, and we found that there was almost no reason to have that link in there. Independent of that Howie Frank's group came in, and they did a design, and they also felt that that link was unnecessary. So you see, that takes some daring, because if you mess around, if you make a mistake you are in trouble. At this point some intelligence was built into the network. Howie always made the point that network flow theory is a really complicated theory and you can't do network flow design on the back of an envelope. If you look at this network, this network looks like the back of an envelope. (Laugh) So NAC and we at UCLA were very friendly, professionally, personally; but we also were competitive, because one of my students, essentially Mario Gerla and myself, were busy doing network design, and developing network design procedures. I had done some, and Mario and I developed some more. And we developed a network design algorithm called CBE, the Concave Branch Elimination method. NAC had theirs, called the Cut Saturation Method. They were competitive, and they were both effective. But theirs was the one that was used because they were under contract. We later showed, through some research, that even though they came up with a very different network, it had about the same cost performance profile as networks we were generating, which is rather interesting. So we were competitive but very cooperative. Mario then went to work for NAC for about three or four years. He worked with Wushow Chou, who is now at North Carolina State University. We then hired Mario back here, after he had seasoned out there. So that was one of our relationships.

Now for BBN. I characterize the BBN relationship this way. BBN, as were most software groups at that time, were largely uninterested in performance. That's an exaggeration. But by and large a programmer, a software guy, simply wants to get a piece of software that works. That's hard enough. Whether it works efficiently or well is not usually the issue; they worry about that maybe way down the line. So BBN was not focusing on efficiency or performance but on making it work and worrying about the logic. So we had a running battle with them in terms of showing them how to improve performance. Another student of mine (Gary Fultt) and I showed them early on that their routing procedure was inferior to two others we developed. We did this in 1970. Finally in 1979 they implemented something very much like one of our procedures and then some years later another piece of it. BBN monitored the network. They measured traffic, and they measured line failures, and IMP failures. We were doing things like measuring response times (which they were not doing) and stressing the network by generating extra network traffic to see what would happen. We would, almost on demand, break the network. There were faults in the software, logic faults. So you can see BBN was not very happy with us showing up their faults and telling them to fix them. Sooner or later they did. Any large system is going to contain faults. But we became experts in finding these things, and I document a whole bunch of them in here. One time we were making measurements, which means we turned on the IMP software which made measurements and which were sent back to us at the Network Measurement Center. As a result of that the network kept crashing. So BBN called us up and said, "You're crashing the network!" And we said, "We need to make measurements." And we found out why. Are you a technical person at all?

O'NEILL: Yes, somewhat.

KLEINROCK: Well, the IMP had only so much storage, initially a very limited amount of storage. And some of it was set aside to reassemble packets into messages - the reassembly function. Specifically, the ARPANET would take a message and break it up into multiple packets. They will not be delivered until all of its packets were collected at the destination. So when packets arrive at their destination, you have to put aside some buffer space to reassemble them. BBN made the calculation and said, "Look, if I've got this many buffers, what is the

maximum number of messages I can reassemble at the same time?" Well, the smallest multi-packet message is two packets. One packet doesn't have to be reassembled. So they took the size of the storage and divided by two to determine how many messages could be in the process of reassembly. And that is how many pointers they installed in their software. A few years later, in 1973, they increased the storage by a factor of two, and they forgot to increase the number of pointers. Who would ever think of that. So they could allocate space, but could never point to it again! CRASH! It turns out that the measurement messages are two packets long. So a pile of them came into our switch, took us down as well as the rest of the network. A similar thing happened just last fall to AT&T. So the UCLA-BBN relationship was one of guarded respect.

O'NEILL: When you had problems, or when you could see that there was a problem with the network, or when you were trying to get something accomplished, did you go through the ARPA office - through Larry Roberts? Or did you deal with BBN directly?

KLEINROCK: We dealt with BBN directly. When we had a problem with BBN, we complained to Larry and he would step in and make sure things were fixed up. It was not a formal relationship that required all kinds of paperwork to go back and forth. It was peers, and researchers, and developers. It was a friendly and efficient environment in that sense.

O'NEILL: I'm trying to get a better feel for how Roberts influenced the development of the ARPANET.

KLEINROCK: He was dominant. In the sense that he set goals, set things in motion, and protected us from all the vagaries of the ARPA structure on the other side of him. He managed to keep my side simple - he was a master at that. Have you met Larry?

O'NEILL: No.

KLEINROCK: Larry is about my age, in his early fifties. He's one of the smartest guys I know. But he wouldn't always fill in all the blanks when he was saying something. So people sometimes had a hard time following him. The things he accomplished while at ARPA were phenomenal!

TAPE 2/SIDE ONE

KLEINROCK: The ARPANET development took place during Larry's tenure as did part of the packet satellite experiment. Somewhere around that time Bob Kahn came in. It may have been at the end of the packet satellite effort. Certainly the packet radio project was pretty much Bob's baby as were many things beyond that. You mentioned SRI. SRI was an important member of the community; as you see, they were the second node in the ARPA network. They served as Network Information Center, NIC. We did not have much of a relationship with them. Documentation went through them, as did the network RFC notes. BBN was the principal player so far as we were concerned in the ARPANET. There's a bit of an anecdote here. It has to do with the arrival of the IMP at UCLA. It was supposed to be a week late because they were slow at BBN - we heaved a sigh of relief because we needed the time. They put the damn thing on an airplane and air freighted it out here - it appeared on time! That next morning everybody was present here. My staff, the computer science chairman, the school of engineering administration, somebody from the chancellor's office, somebody from AT&T long lines, local telephone company people, Honeywell, the DARPA guys, the BBN guys - and everybody was ready to point the accusing finger! Fortunately that first day we got bits moving back and forth properly. The next day messages were moving back and forth. I have to give credit to BBN; they did a superb job. Our end worked well; their end worked well. It went off with no hitches.

O'NEILL: So on the same day that it was delivered...

KLEINROCK: We were moving bits back and forth. We had a guy here, Mike Wingfield, who was in charge of our side of the hardware interface; he was excellent. We have a huge number of photographs with the key

people in them, if you are interested in that; they came back for our ACT ONE symposium. There were a lot of people involved in this project.

O'NEILL: What kind of relationship did you have, if any, with the wider host community?

KLEINROCK: We at UCLA were charged with making the host-to-host software. Steve Crocker was in charge of that effort. It was supposed to be completed by the summer of 1970; it started in the summer of 1969.

O'NEILL: This is the Network Working Group?

KLEINROCK: That's right. A bunch of graduate students were given this job. It was not until October 1971 when the specification of the host-to-host protocol was complete and it began to be installed. It was more than a year late. It was a much harder project than we thought, and it was done not only locally at UCLA, but among various hosts that were involved, including people at the University of Utah and SRI. So the way the ARPANET was used until October 1971 was not easily, but it was used mainly as people migrated from one site to another (typically when they changed jobs). They wanted to use the software back at their old hosts, and they knew how to use it; that generated a lot of early use. The host-to-host interface was awful to begin with. Before the network was installed, I interacted with the early host sites in the following way (as reported in one of my early published papers): We tried to get a feeling for how much traffic would be placed on the network between different nodes. So I called the individual sites and I said, for example, "MIT, how much traffic are you going to send to Illinois?" And the principal investigator said, "I don't know; I have no idea how much I'm going to send." "Will you send, like, three teletypes' worth?" and they said, "Yes, three teletypes' worth." So I collected all this data, and I published a paper with that particular traffic matrix. It had no relationship to reality whatever. There was an interesting issue here that affected the architecture of the ARPANET. When Larry Roberts and Bob Taylor announced this network, most of the hosts did not want to participate. They considered this an impediment. Attachment to the network implied that they would have to implement some host software; that

is, they would have to do major surgery on their operating system to get this network to talk to them. Once the host-to-host protocol was specified that implied they had to implement it on their machine. Now if they had the same kind of machine, like a PDP-10, that someone else had, then fine, but if they had something else like an ILLIAC IV, they had to do it themselves. Even early on, people were not happy about this. They didn't want people using their machine, etc. So Larry had to do a sales job and convince people that there was good reason to implement it. As a result, some of the ways in which the network was specified was done such that it would not impact the user very much. That made it more complicated for the internal operations of the network to do things in the network that should have been done outside the network. In general, people said, "No, I don't want to participate." And we managed to get most of them to participate eventually, as you can see.

O'NEILL: Did they come right out and say, "No, we don't want to do this," or was it just kind of a general feeling?

KLEINROCK: Well, Larry couldn't say, "Thou shalt."

O'NEILL: But he was controlling the money.

KLEINROCK: Yes, but he wasn't that crude. He encouraged them. Larry was not dictatorial that way; he was decisive, but he wouldn't really force people. He would cajole them, strongly. They soon began to realize that there was a benefit. You see the biggest surprise about the ARPA network use was e-mail. It was an ad hoc add-on by BBN, and it just blossomed. And that sucked a lot of people in. It still is the biggest use of networks today. Originally, the network was supposed to provide resource sharing, not even data sharing. For example, you would log on to Utah to use their graphics capability there. At one time it was thought maybe you could import the software to your machine and run it locally. But the original concept was that you would do resource through the network - that's not really what happened. What happened was we used it as a communications medium for access to data, as opposed to access to programs. Early on, Rand was using the

machine here at UCLA. The guys at Rand had earlier been carrying decks of cards over to UCLA to run their weather simulation and climate simulation on our big 360/91 at the time. Once the network came on line, they generated a lot of traffic going back and forth; it was really good stuff, too. So the community of people using the network was the scientists as well as the computer-types who were using computing for computing's sake, and not for computer science's sake.

O'NEILL: Were you using the network early on as well? Were you using the e-mail and seeing these benefits for yourself?

KLEINROCK: E-mail, yes. I wasn't using the external graphics. Neither did I use resource sharing. I used it as a vehicle for studying networks.

O'NEILL: I want to understand what was meant by resource sharing. Was it the same as distributed processing? Was that a goal?

KLEINROCK: That is a concept nowadays. Actually it was and was not. Let me preface that by asking, are you familiar with the October 1972 ICCC public demonstration of the ARPANET?

?? or IEEE

O'NEILL: Yes.

KLEINROCK: It was announced months ahead of time. All of us were told we were going to have this demonstration. DARPA installed an IMP in a hotel in Washington, D.C. and ran in some lines. Everybody was encouraged to create some demonstration packages, and we did as well. That caused lots of good things to happen in the ARPA network. It generated lots of new uses of the ARPA network, just for that demo. One of the things that was demonstrated there was a distributed air traffic control system. The idea was there would be a bunch of computers in the network that would be simulating air traffic control operation in their physical

region. For example, MIT would be doing Boston, and some Washington machines would do Washington and so on in different regions. The idea was there would be some air traffic on one of the computers. As this plane moved out of its region, it would be picked up by another computer's air traffic control, data would be exchanged, and collectively they would be managing this whole air space over a large region. This package also had the ability to freeze the simulation, take the program on machine A, squeeze it over to machine B, and then continue the simulation. There were really some sophisticated things going on there. I remember one of the demos was really interesting. In this demo, you would sit down in Washington at a teletype, logon to a machine at BBN, pull up some source code, ship it over to a machine at UCLA across the country, compile and execute, and bring back the results to be printed on the teletype right next to you in Washington. Now at the same time there were some artificial intelligence demonstrations going on. There was a thing called Turtle running around the ground, bumping into things and avoiding them - MIT had that. Jon Postel sat down to demonstrate this thing: logon to BBN, shift it over to UCLA, compile and execute, and then send it back to print. Nothing happened! He couldn't figure out what was wrong. He kept looking around. Then he found Turtle on the floor, and Turtle was jumping around... It turns out that the output was going to Turtle - they had messed up a connection. There's your output jumping around on the floor. (Laugh) But the point is it was a great demo. People were pulled out of the hallway, handed a handbook, and told, "Sit down, we'll help you use the ARPANET," and they could. So that was another case where a lot of hosts got involved, and distributed processing came about. But that was not the main purpose. The main purpose was to prove networking. And from Larry's point of view as the person funding all of this, he was pleased to get some user sharing. And Larry did a paper early on to show the cost effectiveness of the ARPANET. He said if you had to replicate all of this at every site, it would cost you millions and millions; and here was a network providing this capability at far less cost; therefore we win. This was a false argument. No one was going to reproduce the ILLIAC IV at all of these places.

O'NEILL: You mentioned that Roberts shielded you and the other contractors from a lot of what was going on behind the scenes at ARPA, or that's how you felt. Did you have any exposure to the military aspect of ARPA? Did you ever feel like you had to have concern for DoD interests?

KLEINROCK: Yes. Every time I wrote a proposal I had to show the relevance to the military's applications. The whole packet satellite, and certainly the packet radio, was clearly military - soldiers running around, tanks, troop carriers, aircraft, submarines. Yet, we never did any military work. We realized there was an application there, but that didn't drive our research or our experimentation. It may be that people like Rockwell, who developed the packet radio, were told, "It has to be small and lightweight, and low power" and all the rest. In fact, the survivable aspects of it, I'm sure, were dictated by the military. You could chop pieces out of it and still make it work. But we had only to identify these applications, and to point out how this body of knowledge would be used for that. It was not at all imposed on us. But it definitely colored our proposals. I mean, we were well aware that we had to put that in in order to get funding.

O'NEILL: Was that from the very beginning or did that change? Do you recall much difference?

KLEINROCK: It got worse and worse as time went on. More and more - instead of worse and worse.

O'NEILL: You mentioned that some of the early people who were to be hosts were negative about the idea of sharing their resources. Were you exposed to other negative reactions to the network? For example, people have mentioned the phone company being very negative.

KLEINROCK: Oh, yes, that was clear. It has been said that the telephone industry, or the communications industry, had absolutely nothing to do with the development of the ARPA network, of packet networks. To first order that's correct. You know, IBM was one of the people at that 1967 meeting, as I recall. They sent a fellow named Doug McElroy who was a good man, interested, with good ideas. And then IBM dropped out, and I

think the reason was that they wanted to bid on this thing, and they couldn't be part of the inner circle and also bid. They never appeared again in the community at all. Of course, they went SNA. Also, AT&T was not involved as an organization. I remember in the late 1960s what was going on at the various computer conferences. There would always be a panel consisting of both industries, data communications and data processing. There would be a debate between those two. The computer guys would say, "Communication guys, will you please give us good data communications." The communications guys would turn around and say, "What are you talking about, the United States is a copper mine. You've got wires all over the place; use them." The computer guys would say, "No, you don't understand. It takes half a minute to set up a call, and your charge is for a minimum of three minutes. All I want to do is send a hundred milliseconds of data." These guys would turn around back to the computing guys and say, "Go away little boy, there's no revenue there." So the little boys went away, and they created packet switching. There was no revenue, and in some sense, there still isn't sufficient revenue. The way the communication carriers are going to make money out of data is to digitize voice and to offer new data service. They've done the former by packetizing voice. Now they have made a digital network out of the old analog network, but were reluctant to do so. In fact, just look at the history. They didn't get into networking until GTE bought up Telenet in 1979. In April 1978, AT&T was supposed to develop something called the Bell Data Network. They never did. In 1979 they started talking about AIS, Advanced Information Service, another network they never made. Then finally in 1982, they came out with Net 1000. In 1986 they closed down Net 1000; they lost a billion dollars on that effort. Now they have Accunet and more. It took them decades to come up to the technology that the data processing guys developed in the ARPANET.

O'NEILL: Was the term "packet switching" used all along? There were some early papers where they refer to message switching, although the idea of the packet was there.

KLEINROCK: The idea of the packet was there from the initial days of the ARPANET. It's also implicit in my 1964 book, "Communication Nets; Stochastic Message Flow and Delay."

O'NEILL: Maybe we should wrap this up because I know you have an appointment coming up. First let me ask you just one or two questions about Act One. I did see the people that you had invited. What was the goal of Act One? And did you invite the people that you thought had been the main players or are currently the active players? How did you decide who to invite?

KLEINROCK: There were two reasons for ACT ONE. First, it was the twentieth anniversary of the ARPANET, and the ARPANET we knew was going to be killed. So we felt we should commemorate that event. Second, UCLA did want to start a symposium series, and it was a great way to launch it. That was the motivation. Steve Crocker was one of the people who originated the idea of a commemoration. He made the point, "Look, ARPANET is going away; it's twenty years; let's do something." So he was one of the prime movers. So he, Jerry Estrin, and I were the key people in bringing about the ACT ONE symposium. Jerry and I turned out to be co-chairmen because we are with the university, but Steve was one of the inner members. And we invited the people to each session that we thought were appropriate: the ARPANET history with the ARPANET history guys, today's technology, and the future. We asked various of the players who they thought should be invited. And some people could make it, some people couldn't. We ended up with excellent representation. In fact, in this case, the audience was every bit as qualified as the speakers. We had a lot of the guys out there - the old guys and the new guys. I don't know if you heard, but the venue, the whole flavor, was first rate... Did you know there was a lot of poetry at ACT ONE?

O'NEILL: Alex McKenzie had some of it that I took a look at.

KLEINROCK: It was just a fun thing. It was done with class; it was a class act. I've got the poetry somewhere, because I wrote some of it. Vint wrote some, and Barry Boehm wrote some.

O'NEILL: Do you have any general comments about the ARPANET, or about its development, that you would like to add?

KLEINROCK: Yes. I think that DARPA IPTO (now ISTO) was a prime mover for the United States in the advancement of computer technology through advanced thinking and... what shall I say... heroic funding of the things they thought were worthwhile. Their motto was, "High risk, high payoff." That's exactly what they did with a number of projects. You named them before: time-sharing, networking, AI, a few other areas. Networking was one of their major successes. They backed it fully; lots of money, lots of freedom in terms of what we were doing, really advanced technology. It was a great, great experiment - I can't applaud them more. It was one of the great experiments in science, I think. It completely changed the way things are going on now - commerce, government, industry, science, etc. Who to thank? The key players at DARPA for sure. But I think the originators of the whole ARPA concept - Licklider being one of them. The concept of "Let's fund the kinds of projects and the people that we believe know how to carry out these projects well." And they cut through a lot of red tape. They didn't go through all the peer reviews; they just bet on people that they had confidence in. I think it worked real well. It was a heroic kind of thing. On a scale of one to ten I would give them a ten. Of course, there were some bad things along the way. Whenever you are aggressive and advanced, there is going to be something and somebody left out; you can't help that.

O'NEILL: So that was the price that you paid for the success.

KLEINROCK: It was not done in a democratic way. Not every step was carefully judged, "Should we do this, that, or the other thing? Let's get thirty people to vote." Some guy said, "Look, I believe that is right - let's do it." And they did it. And the money was forthcoming. There were a lot of hurdles, telephone company hurdles, government hurdles, university hurdles - all the way. And that was not only for the ARPANET project but also other projects as well.

O'NEILL: You mentioned that Roberts was shielding you, or that you felt he was shielding you, from what was going on behind the scenes. Do you know what was going on? What he was doing? Or did you just see the end result of that - that you didn't have to hassle with it?

KLEINROCK: Now for BBN. I characterize the BBN relationship this way. **BBN, as were most software groups at that time, were largely uninterested in performance.** That's an exaggeration. But by and large a programmer, a software guy, simply wants to get a piece of software that works. That's hard enough. Whether it works efficiently or well is not usually the issue; they worry about that maybe way down the line. So BBN was not focusing on efficiency or performance but on making it work and worrying about the logic. So we had a running battle with them in terms of showing them how to improve performance. Another student of mine (Gary Fultt) and I showed them early on that their routing procedure was inferior to two others we developed. We did this in 1970. Finally in 1979 they implemented something very much like one of our procedures and then some years later another piece of it. **BBN monitored the network. They measured traffic, and they measured line failures, and IMP failures. We were doing things like measuring response times (which they were not doing) and stressing the network by generating extra network traffic to see what would happen. We would, almost on demand, break the network. There were faults in the software, logic faults. So you can see BBN was not very happy with us showing up their faults and telling them to fix them**. Sooner or later they did. Any large system is going to contain faults. But we became experts in finding these things, and I document a whole bunch of them in here. One time we were making measurements, which means we turned on the IMP software which made measurements and which were sent back to us at the Network Measurement Center. As a result of that the network kept crashing. So BBN called us up and said, "You're crashing the network!" And we said, "We need to make measurements." And we found out why. Are you a technical person at all?

O'NEILL: Yes, somewhat.

KLEINROCK: Well, the IMP had only so much storage, initially a very limited amount of storage. And some of it was set aside to reassemble packets into messages - the reassembly function. Specifically, the ARPANET would take a message and break it up into multiple packets. They will not be delivered until all of its packets were collected at the destination. So when packets arrive at their destination, you have to put aside some buffer space to reassemble them. BBN made the calculation and said, "Look, if I've got this many buffers, what is the maximum number of messages I can reassemble at the same time?" Well, the smallest multi-packet message is two packets. One packet doesn't have to be reassembled. So they took the size of the storage and divided by two to determine how many messages could be in the process of reassembly. And that is how many pointers they installed in their software. A few years later, in 1973, they increased the storage by a factor of two, and they forgot to increase the number of pointers. Who would ever think of that. So they could allocate space, but could never point to it again! CRASH! It turns out that the measurement messages are two packets long. So a pile of them came into our switch, took us down as well as the rest of the network. A similar thing happened just last fall to AT&T. So the UCLA-BBN relationship was one of guarded respect.

X-Sender: walden@labs-n.bbn.com
Date: Fri, 03 Feb 1995 09:12:00 -0600
To: katieh@bga.com
Subject: comments on len's comments

<<I put my comments on Len's comments in double angle brackets.

You should also check what I said about bugs in the IMP in my interview by Judy O'Neil.>>

Now for BBN. I characterize the BBN relationship this way. BBN, as were most software groups at that time, were largely uninterested in performance.

<<BBN was not just a software group. BBN was strong in hardward, software, theory and systems design and implementation. For instance, Ornstein, Barker and Heart were hardware people. Crowther Cosell and I were software people. Kahn and Crowther were powerful with theory. All of us had successfully built complex systems. Our greatest strength then and for years to come was the ability to
There are 141 lines left (16%). Press <space> for more, or 'i' to return.

Our greatest strength then and for years to come was the ability to find the best mix between hardware and software.>>

That's an exaggeration. But by and large a programmer, a software guy, simply wants to get a piece of software that works. That's hard enough. Whether it works efficiently or well is not usually the issue; they worry about that maybe way down the line. So BBN was not focusing on efficiency or performance but on making it work and worrying about the logic.

<<It appears to me that one needs to define what one means by "performance." When Len says "performance," he appears to be talking about the efficiency and correctness network wide algorithms. I'll come back to these in a minute, but we were also very concerned with "performance" in the sense of fitting a lot of complexity into a little box and making the little box run very very fast. We were also concerned with efficiency of network wide algorithms, and this actually was at odds with us wanting to make some of the changes that would lead to more correct network wide algorithms. Many software groups have the characteristics Len describes, but our 3 person group
There are 121 lines left (28%). Press <space> for more, or 'i' to return.

the characteristics Len describes, but our 3 person group was enormously focused on program (it not always network) performance, and in fact we had calculated out the program performance in advance of writing it.

Also, BBN didn't speak with one voice about the issues Len brings up. For instance, you should get Kahn's views on the situation. He views himself as having been in a running battle with Crowther about the correctness of the network algorithms. Also, farily early, before the work by Len's students, I think, Kahn and I ran the famous tests at UCLA that showed some of the problems with the initial network algorithms and which Kahn views as having proved that Crowther should have listened to him better.

After that Kahn and Crowther wrote a report describing
the problems.>>

So we had a running battle with
them in terms of showing them how to improve performance. Another
student of mine (Gary Fultt) and I showed them early on that their
routing procedure was inferior to two others we developed. We did
There are 101 lines left (40%). Press <space> for more, or 'i' to return.
Message 1/93 From David C. Walden Page 4

routing procedure was inferior to two others we developed. We did
this in 1970. Finally in 1979 they implemented something very much
like one of our procedures and then some years later another piece of
it. BBN monitored the network. They measured traffic, and they
measured line failures, and IMP failures. We were doing things like
measuring response times (which they were not doing) and stressing the
network by generating extra network traffic to see what would happen.
We would, almost on demand, break the network. There were faults in
the software, logic faults. So you can see BBN was not very happy
with us showing up their faults and telling them to fix them. Sooner
or later they did. Any large system is going to contain faults. But
we became experts in finding these things, and I document a whole
bunch of them in here. One time we were making measurements, which
means we turned on the IMP software which made measurements and which
were sent back to us at the Network Measurement Center. As a result
of that the network kept crashing. So BBN called us up and said,
"You're crashing the network!" And we said, "We need to make
measurements." And we found out why. Are you a technical person at
all?

<<I don't view there as having been much of a running battle,
There are 81 lines left (52%). Press <space> for more, or 'i' to return.
Message 1/93 From David C. Walden Page 5

<<I don't view there as having been much of a running battle,
and I was responsible for operation of the network for much
of this time. In fact, from the earliest days we were also
writing papers and reports on problems we discovered or fixes
to problems that others discovered. There could have been
some tension between our job of keeping the network up and running
and them scheduling time to test it and possibly break it,
but I always felt that that was mostly worked out easily enough.

As Len describes, they had the job from ARPA of doing network
measurement, and we mostly spent our time doing the job that we
had from ARPA, which was a different job. We got their reports,
tried to figure out ways to correct the problems they discovered,
we used some of their suggested methods and didn't use some
others. My view is that it was all pretty collaborative and
together we did some pretty amazing stuff.

I do believe to this day that some of the work by Len and
his people was overly theoretical, i.e., ignored too many
real life issues, and that my experience with Len
was one of the things that made me really happy to
There are 61 lines left (63%). Press <space> for more, or 'i' to return.
Message 1/93 From David C. Walden Page 6

was one of the things that made me really happy to
be in industry where people care less about who gets credit
for what is done and more about getting something done. Seeing

the academic inclination hold back telling us about things until they were published, when we would tell them anything as soon as we knew it was a real eye-opener to me. Of course, Len may see it exactly in reverse.

As I said above, in general, I have the utmost respect for the work Len and his people did, and think that our implementation inclination combined with their empirical inclination produced as better result that either of us would have produced alone.

Speaking for me personally, I never felt very involved in any of the tension Len describes. Mostly I just bopped along telling anyone who would listen what we were doing, listening to whatever anyone else said, and trying to make it all better as best we could.>>

O'NEILL: Yes, somewhat.

There are 41 lines left (75%). Press <space> for more, or 'i' to return.
Message 1/93 From David C. Walden Page 7

KLEINROCK: Well, the IMP had only so much storage, initially a very limited amount of
storage. And some of it was set aside to reassemble packets into messages - the reassembly function. Specifically, the ARPANET would take a message and break it up into multiple packets. They will not be delivered until all of its packets were collected at the destination. So when packets arrive at their destination, you have to put aside some buffer space to reassemble them. BBN made the calculation and said, "Look, if I've got this many buffers, what is the maximum number of messages I can reassemble at the same time?" Well, the smallest multi-packet message is two packets. One packet doesn't have to be reassembled. So they took the size of the storage and divided by two to determine how many messages could be in the process of reassembly. And that is how many pointers they installed in their software. A few years later, in 1973, they increased the storage by a factor of two, and they forgot to increase the number of pointers. Who would ever think of that. So they could allocate space, but could never point to it again! CRASH! It turns out that the measurement messages are two packets long. So a pile of them came into our switch, took us down as well as the rest of the network. A
There are 21 lines left (87%). Press <space> for more, or 'i' to return.
Message 1/93 From David C. Walden Page 8

into our switch, took us down as well as the rest of the network. A similar thing happened just last fall to AT&T. So the UCLA-BBN relationship was one of guarded respect.

<<I don't actually remember the instance Len describes, but such bugs do get put into programs. Being really experienced implementors, being caught with bugs wasn't very threatening to us. It was more of a "damn, made another big mistake in public again, big sigh, well, how do we fix this one."

My respect for UCLA was never guarded. I was also enormously impressed with Kleinrock, Cerf, Crocker, Postel, Wingfield, Estrin, Naylor, Fultz, Lam, Gerla, etc. etc. etc. and I loved having the opportunity to work with them.>>